Researchers improve artificial intelligence algorithms for semi-autonomous vehicles

February 3 2015, by Aaron Dubrow

A vision for urban transportation in 2030. GM-Segway EN-V Laugh (Xiao) exhibited at the Shanghai Expo 2010. Credit: Wikimedia Commons

For decades, researchers in artificial intelligence, or AI, worked on specialized problems, developing theoretical concepts and workable

algorithms for various aspects of the field. Computer vision, planning and reasoning experts all struggled independently in areas that many thought would be easy to solve, but which proved incredibly difficult.

However, in recent years, as the individual aspects of artificial intelligence matured, researchers began bringing the pieces together, leading to amazing displays of high-level intelligence: from IBM's Watson to the recent poker playing champion to the ability of AI to recognize cats on the internet.

These advances were on display this week at the 29th conference of the Association for the Advancement of Artificial Intelligence (AAAI) in Austin, Texas, where interdisciplinary and applied research were prevalent, according to Shlomo Zilberstein, the conference committee chair and co-author on three papers at the conference.

Zilberstein studies the way artificial agents plan their future actions, particularly when working semi-autonomously—that is to say in conjunction with people or other devices.

Examples of semi-autonomous systems include co-robots working with humans in manufacturing, search-and-rescue robots that can be managed by humans working remotely and "driverless" cars. It is the latter topic that has particularly piqued Zilberstein's interest in recent years.

The marketing campaigns of leading auto manufacturers have presented a vision of the future where the passenger (formerly known as the driver) can check his or her email, chat with friends or even sleep while shuttling between home and the office. Some prototype vehicles included seats that swivel back to create an interior living room, or as in the case of Google's driverless car, a design with no steering wheel or brakes.

Except in rare cases, it's not clear to Zilberstein that this vision for the vehicles of the near future is a realistic one.

"In many areas, there are lots of barriers to full autonomy," Zilberstein said. "These barriers are not only technological, but also relate to legal and ethical issues and economic concerns."

In his talk at the "Blue Sky" session at AAAI, Zilberstein argued that in many areas, including driving, we will go through a long period where humans act as co-pilots or supervisors, passing off responsibility to the vehicle when possible and taking the wheel when the driving gets tricky, before the technology reaches full autonomy (if it ever does).

In such a scenario, the car would need to communicate with drivers to alert them when they need to take over control. In cases where the driver is non-responsive, the car must be able to autonomously make the decision to safely move to the side of the road and stop.

The driving routes (in green) computed by a Lexicographic Value Iteration (LVI) algorithm for an attentive driver (above) and a tired driver (below) based on traffic and road conditions. Credit: Shlomo Zilberstein, University of Massachusetts Amherst

"People are unpredictable. What happens if the person is not doing what they're asked or expected to do, and the car is moving at sixty miles per hour?" Zilberstein asked. "This requires 'fault-tolerant planning.' It's the kind of planning that can handle a certain number of deviations or errors by the person who is asked to execute the plan."

With support from the National Science Foundation (NSF), Zilberstein

has been exploring these and other practical questions related to the possibility of artificial agents that act among us.

Zilberstein, a professor of computer science at the University of Massachusetts Amherst, works with human studies experts from academia and industry to help uncover the subtle elements of human behavior that one would need to take into account when preparing a robot to work semi-autonomously. He then translates those ideas into computer programs that let a robot or autonomous vehicle plan its actions—and create a plan B in case of an emergency.

There are a lot of subtle cues that go into safe driving. Take for example a four-way stop. Officially, the first car to the crosswalk goes first, but in actuality, people watch each other to see if and when to make their move.

"There is a slight negotiation going on without talking," Zilberstein explained. "It's communicating by your action such as eye contact, the wave of a hand, or the slight revving of an engine."

In trials, autonomous vehicles often sit paralyzed at such stops, unable to safely read the cues of the other drivers on the road. This "undecidedness" is a big problem for robots. A recent paper by Alan Winfield of Bristol Robotics Laboratory in the UK showed how robots, when faced with a difficult decision, will often process for such a long period of time as to miss the opportunity to act. Zilberstein's systems are designed to remedy this problem.

"With some careful separation of objectives, planning algorithms could address one of the key problems of maintaining 'live state', even when goal reachability relies on timely human interventions," he concluded.

The ability to tailor one's trip based on human-centered factors—like

how attentive the driver can be or the driver's desire to avoid highways—is another aspect of semi-autonomous driving that Zilberstein is exploring.

Google self-driving car in Mountain View, California. Credit: Mark Doliner via Wikimedia Commons

In a paper with Kyle Wray from the University of Massachusetts Amherst and Abdel-Illah Mouaddib from the University of Caen in France, Zilberstein introduced a new model and planning algorithm that allows semi-autonomous systems to make sequential decisions in situations that involve multiple objectives—for example, balancing safety and speed.

Their experiment focused on a semi-autonomous driving scenario where the decision to transfer control depended on the driver's level of fatigue. They showed that using their new algorithm a vehicle was able to favor roads where the vehicle can drive autonomously when the driver is fatigued, thus maximizing driver safety.

"In real life, people often try to optimize several competing objectives," Zilberstein said. "This planning algorithm can do that very quickly when the objectives are prioritized. For example, the highest priority may be to minimize driving time and a lower priority objective may be to minimize driving effort. Ultimately, we want to learn how to balance such competing objectives for each driver based on observed driving patterns."

It's an exciting time for artificial intelligence. The fruits of many decades of labor are finally being deployed in real systems and machine learning is being adopted widely and for different purposes than anyone had ever realized.

"We are beginning to see these kinds of remarkable successes that integrate decades-long research efforts in a variety of AI topics," said Héctor Muñoz-Avila, program director in NSF's Robust Intelligence cluster.

Indeed, over many decades, NSF's Robust Intelligence program has supported foundational research in artificial intelligence that, according to Zilberstein, has given rise to the amazing smart systems that are beginning to transform our world. But the agency has also supported researchers like Zilberstein who ask tough questions about emerging technologies.

"When we talk about autonomy, there are legal issues, technological issues and a lot of open questions," he said. "Personally, I think that NSF

has been able to identify these as important questions and has been willing to put money into them. And this gives the U.S. a big advantage."

Provided by National Science Foundation